NPS-OR-92-015

NAVAL POSTGRADUATE SCHOOL

Monterey, California

[i.e. PA[illegible]m] PACISM: A Design Aid for the Development of the Petite Amateur Navy Satellite

Russell Gottfried
Michael Bailey

2 September 1992

Prepared for: Naval Postgraduate School
Monterey, CA 93943

**NAVAL POSTGRADUATE SCHOOL,
MONTEREY, CALIFORNIA**

Rear Admiral R.W. West Jr.
Superintendent

Harrison Shull
Provost

This report was prepared for and funded by the Naval Postgraduate School.

This report was prepared by:

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a REPORT SECURITY CLASSIFICATION: UNCLASSIFIED	1b RESTRICTIVE MARKINGS
2a SECURITY CLASSIFICATION AUTHORITY	3 DISTRIBUTION/AVAILABILITY OF REPORT: Approved for public release; distribution is unlimited.
2b DECLASSIFICATION/DOWNGRADING SCHEDULE	
4 PERFORMING ORGANIZATION REPORT NUMBER(S): NPSOR-92-015	5 MONITORING ORGANIZATION REPORT NUMBER(S)
6a NAME OF PERFORMING ORGANIZATION: Naval Postgraduate School; 6b OFFICE SYMBOL (If applicable): OR	7a NAME OF MONITORING ORGANIZATION: Naval Postgraduate School
6c ADDRESS (City, State, and ZIP Code): Monterey, CA 93943	7b ADDRESS (City, State, and ZIP Code): Monterey, CA 93943
8a NAME OF FUNDING/SPONSORING ORGANIZATION: NPS Foundation; 8b OFFICE SYMBOL (If applicable)	9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER: OM&N Direct Funding
8c ADDRESS (City, State, and ZIP Code): Naval Postgraduate School Monterey, CA 93943-5000	10 SOURCE OF FUNDING NUMBERS: PROGRAM ELEMENT NO; PROJECT NO; TASK NO; WORK UNIT ACCESSION NO

11 TITLE (Include Security Classification)
PACISM: A Design for the Development of the Petite Amateur Navy Satellite

12 PERSONAL AUTHOR(S)
Michael Bailey & Russell Gotfried

13a TYPE OF REPORT	13b TIME COVERED	14 DATE OF REPORT (Year Month Day)	15 PAGE COUNT
Technical	FROM ______ TO ______	Sept 02, 1992	1

16 SUPPLEMENTARY NOTATION

17 COSATI CODES			18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	

19 ABSTRACT (Continue on reverse if necessary and identify by block number)

In this work, we describe the features and use of PACISM, a simulation model used to support satellite design decisions during the construction of PANSAT, the Petite Amateur Navy Satellite, at the Naval Postgraduate School. The model features communications modeling at the packet level, an accurate portrayal of the store-and-forward mechanism the satellite employs, satellite footprint movement over the earth's surface, and bit error generation, detection, and correction. We describe the object-oriented design of the model, as well as how it is being used by engineers to support decisions about the equipment used on the satellite, the way the equipment will be employed, the way we control the user population, and methods to keep control of the satellite.

20 DISTRIBUTION/AVAILABILITY OF ABSTRACT: UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☐ DTIC USERS	21 ABSTRACT SECURITY CLASSIFICATION: UNCLASSIFIED
22a NAME OF RESPONSIBLE INDIVIDUAL: Professor Michael Bailey	22b TELEPHONE (Include Area Code): (408) 646-2085; 22c OFFICE SYMBOL: OR/Ba

DD FORM 1473, 84 MAR — 83 APR edition may be used until exhausted. All other editions are obsolete

SECURITY CLASSIFICATION OF THIS PAGE
UNCLASSIFIED

PACSIM: A Design Aid for the Development of the Petite Amateur Navy Satellite

Russell Gottfried
Michael Bailey
Department of Operations Research
Naval Postgraduate School
Monterey, CA 93943-5000, U.S.A.

Abstract

In this work, we describe the features and use of PACSIM, a simulation model used to support satellite design decisions during the construction of PANSAT, the Petite Amateur Navy Satellite, at the Naval Postgraduate School. The model features communications modeling at the packet level, an accurate portrayal of the store-and-forward mechanism the satellite employs, satellite footprint movement over the earth's surface, and bit error generation, detection, and correction. We describe the object-oriented design of the model, as well as how it is being used by engineers to support decisions about the equipment used in the satellite, the way the equipment will be employed, the way we control the user population, and methods to keep control of the satellite.

1 NETWORK DESCRIPTION

The Petite Amateur Navy Satellite (PANSAT) is a small, low-orbit spread spectrum communications satellite scheduled for launch in 1995. It is designed for use by amateur radio operators. The spacecraft's orbit is expected to be at an altitude between 450-800 kilometers, at an inclination of 93.5 degrees and a maximum slant range of approximately 2000 km. Users on the ground will be provided with a view of the satellite between four and ten minutes in duration. The size of the user base is currently unknown. However, the fields of spread spectrum, packet radio and satellite communications are widely used among amateur radio operators and attention to this program has grown.

1.1 Summary of Session

The spacecraft will normally be in receive mode until it arrives in the view of a user's ground station and acquires the pseudo-noise coded signal from it. This sequence is a 128 bit stream and is transmitted by the subscriber until acquisition is achieved. The receiver assesses if it is following the sequence in synchronization with the transmitted signal. If not, it slides and looks again to determine whether it matches.

The communications payload of PANSAT is designed for a central frequency of 437.25 Mhz (960 Khz bandwidth). Uplink and downlink will be done within this bandwidth, with information relayed in bit packets. These packets are stored and retrieved in the spacecraft control unit's (SCU) on-board processor. Data storage and retrieval are accomplished using the operators' callsigns as references. PANSAT is intended to have its own address for use by the ground control station. Commands for the satellite are envisioned only to request experiment data and provide telemetry and performance information to the control station.

1.2 High-Level Network Description

The system consists of uncoordinated users who contend for access to a single channel. As shown in Figure 1 these users, also called subscribers, are located at ground sites not necessarily covered by a single orbit of the spacecraft. Due to the periodicity of PANSAT's path, it will not be in a user's view during each orbit.

As the spacecraft comes over the horizon, a subscriber will attempt access. The channel is available for finite communications windows due to the low orbit of the satellite. Users access the channel by locking-up, or synchronizing, with the satellite. When two or more subscribers simultaneously attempt to initiate a session with the spacecraft, there are two possible results: either there is a collision or one will capture the server. A collision results in unsuccessful attempts for both users and requires later scheduled synchronization transmission. Tagged to the end of the synchronization is a preamble identifying the user, followed by the data packet along with

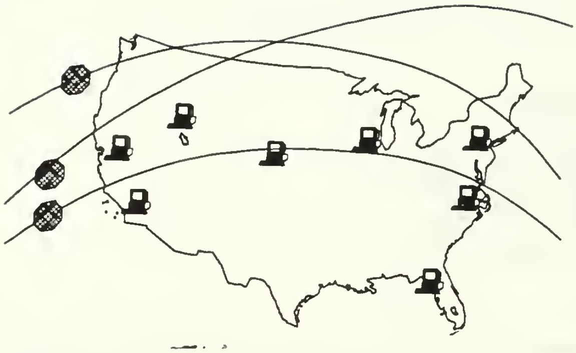

Figure 1: As a low-orbit spacecraft, PANSAT will not be continuously viewed by its users.

its routing instructions, called the header.

If the message arrives at the spacecraft, the SCU

- routes the message for the appropriate addressees;
- forms a queue of messages stored for the addressee;
- transmits, or downloads, the stored traffic after synchronizing with the recipient.

The caller then requests a disconnect and the satellite returns to a ready state, awaiting a new attempt for synchronization.

1.2.1 Channel Access

As mentioned above, users are not coordinated by a master clock nor can they sense whether a channel is being accessed or in use. The only information provided is the starting time of the communications window. This window of time reflects the interval during which the spacecraft is in the view of a user. In other words, the low-earth orbit of the satellite provides a limited horizon on the ground. As this horizon, or footprint, traverses the ground, it covers the position occupied by a user. The period from initial coverage in the footprint to termination is the communications window.

Due to the absence of overall network control, the occurrence of transmission delay and an array of other arbitrary factors, it is assumed that no two callers' synchronizing signals are initiated at the exact same time. However, a characteristic of this transmission-driven protocol is that the synchronization sequence – the elements of which are called chips – is sent iteratively until the SCU's receiver achieves a lock and is captured.

If another attempt occurs such that its sequence is within a phase difference not exceeding the duration of a single chip, the SCU's receiver will be unable to distinguish between the signals and a collision results (Pursely, 1987, p. 118). A collision causes a failed attempt and another is made after a delay. If there is sufficient offset between two synchronization streams, then one will be successful, depending upon which bit stream is more proximal to the pattern expected at the spacecraft's receiver. A caller is made aware of an unsuccessful synchronization if during the time following the signal, no acknowledgement or subsequent message transmission is received from the satellite.

1.2.2 Packet Flow

The message generated by the user is subject to bit error, associated with the bit error rate identified in the spacecraft designers' link budget analysis. Occurrences of bit error, which are not independent events, may signify the loss or partial destruction of a packet. This is unknown to the engaged user. After accessing the channel, the user will accept a synchronization signal from the spacecraft communicating its acknowledgement of receipt of the user's packet and transmitting traffic addressed to the engaged user. If this is not detected by the user, a new transmission is generated in order to ensure receipt of the apparently lost traffic.

After the subscriber transmits an identification preamble, the spacecraft attempts synchronization. There is no contention for the current user's receiver and the sequence is followed by all traffic stored by the spacecraft control unit. Finally, the subscriber follows with an acknowledgment of receipt of the down-linked traffic and a request for disconnect. The SCU recognizes this and returns to a ready state. If the disconnect request is not received, then the SCU times out and retransmits the stored traffic. If final acknowledgment is still not received after a specified period and a stipulated number of retries, the spacecraft returns to a ready state.

1.2.3 Packet Storage and Forwarding

This aspect of message transfer occurs in the spacecraft. Upon arrival of a packet at the spacecraft, it is stored until the SCU is accessed by the addressee. Several factors bear much relevance:

1. The delay in message forwarding from originator to addressee may range from seconds to days.

2. Message size is of arbitrary length although the maximum length may be fixed (Brachman [1]).

3. Packet storage capacity of the SCU is finite.

4. Messages may be addressed to other individual subscribers, multiple users or all users.

Due to storage constraints, as the number of users increases, storage may become scarce. For example, if a packet is addressed to a group of users, it might be beneficial to store single copies of messages with more explicit routing instructions rather than a single copy for each.

2 Network Design Details

This network has associated with it several design issues. Channel access is influenced by the number of users with the spacecraft in view concurrently attempting to access the net, the possibility of collision, the distribution of the duration of the synchronizing procedure as well as the occurrence of bit error and resulting retransmissions of data. Data transfer is influenced by packet length, the occurrence of bit error, propagation delay, the extent of error detection and correction as well as the speed of the SCU's processor. The store-and-forward problem is affected by the number of subscribers, packet routing instructions, bit packet length, and storage capacity. An enumeration of specific design questions and their associated measures of effectiveness is provided in the next chapter.

2.1 Channel Access

Given the absence of coordination among users, it is reasonable to assume that each will independently initiate synchronization attempts in accordance with some arrival process. For a situation in which a geostationary satellite is involved, users have continual access to the spacecraft. In this case, the model may achieve an equilibrium in which attempts to establish communications may be well represented by a carefully selected arrival process. However, because many users are aware of the time at which the spacecraft comes into view, the first attempts by each may be governed by a different timing algorithm than subsequent bids, depending on the occurrence of a success. After arriving, the synchronization process takes place.

When in the ready state, the SCU's receiver is continuously scanning the spreading sequence for the transmission of a synchronization signal. This is a

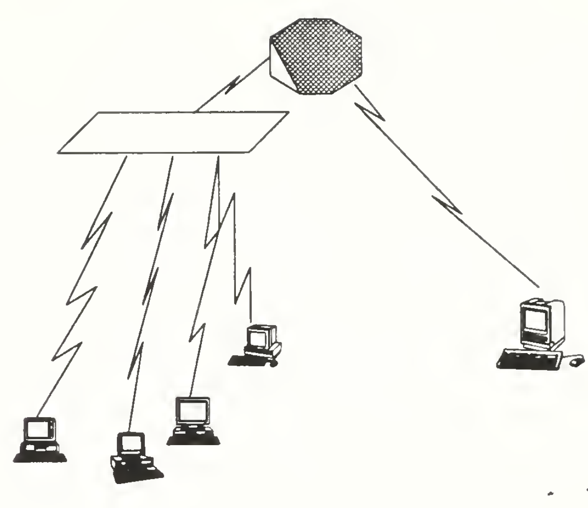

Figure 2: Users contend for SCU access, but there is no contention on the downlink channel

128 bit sequence which is transmitted until an attempt is sensed. At this point, the receiver compares the received pseudo-noise sequence with the expected progression as seen through the current sequence. If there is no match, then the receiver will shift its scanning of the band by a predetermined duration of time and compare again. This is continued iteratively until the receiver attains a lock. The characteristics of the clocking sequence of the pseudo-noise generator dictate the specific quantity of time elapsed during this process.

Between initiation of the synchronizing process and the capture of the channel, there is a non-zero probability of collision. The occurrence of a collision is a function of the distribution of subscribers initiating simultaneous attempts to access the SCU and the distribution of the number of repetitions of the synchronizing sequence which must be completed to achieve lock (Woerner [11]). Regardless of the collision event, multiple users attempting to access the net increases the effective noise in the RF environment. Increased noise causes a deterioration in the link margin and may give rise to increased probability of bit error in session transmissions.

2.2 Data Transfer

This feature of the network is affected by decisions regarding packet size, error detection and correction,

acknowledging message receipt and the establishment of a half- or full-duplex channel.

Decisions concerning packet length have ramifications throughout the operation of the network. Regarding packet transmissions from a subscriber or from the spacecraft, not only do longer packets require a greater amount of time for transmission but they have a greater probability of experiencing bit error. Furthermore, the length of packets impacts the SCU storage capacity. Options include establishing fixed message lengths, maximum message lengths with variable length packets, or leaving message size unconstrained. Each of these considerations may be viewed in light of the bit error rate. Because the occurrence of bit error is not independent among individual bits, the greater the size of the bit packet, the higher the probability of the packet incurring bit error.

The worst-case scenario is that any incident of bit error results in the loss of a packet. However, unless some error detection is instituted, the absence of error in the message header is satisfactory for a message to be received at the destination. This is insufficient for reliable networks (Tanenbaum [10]). To institute negative acknowledgements requires retransmissions and lengthens the span of the session. However, introduction of error correction such as Hamming codes lengthens the packet to more than three times its original length (Tanenbaum [10]), and would also be very likely to have a significant impact on the duration of a session.

The final aspect to be discussed in analyzing the network data transfer is the implementation of a full-duplex as opposed to a half-duplex channel. Regardless which channel type, subscribers independently attempt to synchronize with the spacecraft's receiver. Only if successful will a user be able to conduct a session with the SCU. The session differs conditioned on whether or not the channel is full- or half-duplex. Figures 3 and 4 depict the flow of communications in these two types of curcuits.

In full-duplex, the subscriber's receiver must be in lock with the spacecraft's transmitter in order to be able to conduct a session. Once both caller and SCU are in synchronization, communicated by the user as a connection request and by the spacecraft as an acknowledgement, data exchange occurs. While receiving, acknowledging and storing the data transmitted by the active user, the spacecraft retrieves and forwards messages addressed to the active user. If the packets are error-free and successfully received at either end, then acknowledgements are dispatched and a disconnect occurs. However, if errors are detected then negative acknowledgements are sent and packets

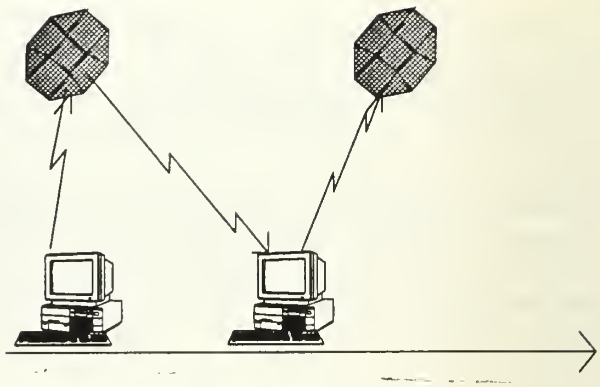

Figure 3: Flow of communications in a full-duplex curcuit.

are retransmitted. Finally, if no acknowledgements are received after a time-out, then the packet is retransmitted.

The half-duplex session is more involved because after each transmission, communications are stopped and the other site must synchronize in order to transmit acknowledgments or data. Specifically, after capturing the SCU and completing the synchronization signal, the active user transmits the data packet.

Upon broadcasting the packets, the user ceases transmission and waits a period of time commensurate with the transmission duration and turn-around time. The SCU is expected to attempt synchronization and transmit an acknowledgement and any mail appropriately addressed. In the case of error-free receipt of the stored traffic, the user again synchronizes with the spacecraft's receiver in order to acknowledge the mail and request disconnect. Once more, the detection of bit error or occurrence of time-out will necessitate retransmission; however, due to the nature of half-duplex communications, retransmissions must be preceded by a synchronization stream.

During the course of these synchronization attempts with the SCU's receiver, the spacecraft will need to prevent intervening callers' attempts to send traffic while awaiting the active user's synchronization to resume the session. Store-Forward

The receipt of packets by the SCU initiates the process of store-and-forward. After the message arrives, it is routed to storage according to the callsign of the addressee. Packets currently in storage and addressed to the current user are retrieved and transmitted.

The network will use a source routing procedure. There are three general cases involving packet routing: single addressee, multiple addressees or mes-

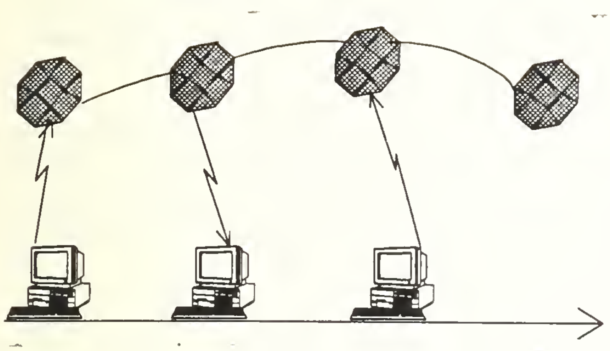

Figure 4: Flow of communications during a half-duplex session.

sages addressed to all users. For single-subscriber addressed packets, once a message has been delivered, processor storage space is made available. For messages addressed to multiple users, a separate listing must be maintained to mark deliveries in accordance with packet headers. These messages may be held for long periods of time.

Regarding the storage constraint, as available memory dwindles, an auto-dump capability ought to be implemented, freeing space for new messages. For effective storage management, packets need to be prioritized in some way, duration of time in storage for example, so that some proportion of the total number of messages are deleted. This brings to light the matter of network accountability in an acknowledged connectionless service. Once the SCU accepts a message, it becomes responsible for delivery of that message. Packets which are dumped due to storage limitations restricts the level of network reliability.

Although straightforward, this communications network has very densely interrelated operating charcteristics. A challenge in modeling and simulating this system is to identify design considerations which may limit operation or which may be altered to enable improved performance. An enumeration of the items of interest and a discussion of the model follows.

3 ENUMERATION OF DESIGN ISSUES

In attempting to create a network which performs well across the areas of channel access, communications flow and packet handling, several questions of interest have emerged. Designers of PANSAT expect that decisions in building the system show significant improvement of performance measures across a broad array of scenarios. Particularly, decision makers seek enhancement in channel access, decrease in session duration and firm assessment of SCU memory requirements. These concerns are enumerated by the following issues and partial listing of applicable design considerations:

- probability of channel access as a function of maximum synchronization sequence length, data transfer rate, the number and distribution of waiting times before retransmission attempts;
- session duration as a function of maximum synchronization sequence length, data transfer rate and maximum (or fixed) packet length;
- SCU storage requirements as a function of maximum packet length and user population density.

As may be determined from this list, there is a strong interdependence among the design issues and operational considerations. In fact, the three areas of activity, access, flow and store-forward, affect one another. Channel access can be expected to be facilitated by shorter sessions. In turn, higher rates of channel access may cut down on variability of the maximum amount of SCU memory in use. To establish the model as a credible decision aid, these issues must be formally defined, analyzed and simulated. This chapter lays out decision makers' questions of interest and measures used in making the determination.

3.1 Channel Access

For PANSAT's network to be successful, users must be furnished with reasonable opportunity to capture the server. This may be measured by the proportion of users obtaining a lock during the first attempt, and during subsequent retries. Capture is affected by:

- maximum number of repetitions of the synchronization sequence;
- data transfer rate;
- discipline governing retries.

The channel access problem may be interpreted as follows. Given a number of callers who simultaneously have the satellite in view, how many attempts, on average, are required to be able to conduct successfully a session with the spacecraft. Alternately, what proportion of users are successful on the first attempt and then on each subsequent attempt.

Underlying the whole channel access question is the concern that the system operator, the Naval Postgraduate School, be assured access to the server. Network design decisions may be implemented to enhance that prospect, but may be found inadequate under the best circumstances. In this case, other means should be developed for activation.

3.1.1 Synchronization Discipline

There is no feedback marking the achievement of lock by any transmitter's spreading code. The user receives no indication of success in capturing the server's receiver until at least after transmitting the identification preamble following the synchronizing sequence. Given an idle server, capture after the full 128 iterations of the spreading code is almost assured; as will be shown, the SCU is captured, on average, in 64 iterations. A lower cap not only reduces the excess time spent transmitting the spreading sequence, but in the event of a busy server, also allows for fewer delays in initiating subsequent retries.

3.1.2 Bit Rate

It is obvious that an increase in data transfer rate will most likely shorten the duration of a session; doubling the bit rate may halve session lengths. Designers' interest in this relationship between bit rate and channel access reflects the notion that shortening the duration of a session will increase users' access to the server. Other factors are not specifically under the engineers' control. These include the effect of the number of collocated users concurrently desiring access or the ramifications of increased bit error rate. As determined by the engineers' communications network link analysis, two feasible data transfer rates are 1200 and 2400 bits per second (bps). However, either the expected bit error rate may increase or the link margin may decrease for the higher data transmission rate (Morgan [7]). The question is whether the higher bit rate significantly improves the likelihood of a user accessing the server despite increased probability of bit error.

3.1.3 Retransmission Regimen

After transmitting the synchronization signal, identification preamble and bit packet, the subscriber is in either of two states. The user

1. is engaged in a session with the spacecraft control unit if the server is idle;
2. will need to retransmit the entire sequence if the server is busy.

In the event of retransmission, what should govern the amount of time to wait and the number of retries? From a network management perspective, the goal is to allow for the largest number of users to access the channel. This might encourage continual retransmissions until success is obtained. On the other hand, it is important to preserve some circuit discipline and minimize noise contributed by other users' transmissions for access to the channel. Resolution of this question will be implemented in a protocol allowing for the highest access rate coupled with the minimum number of attempts.

3.2 Session Duration

In the communications flow paradigm, the spacecraft spends a significant proportion of time in a captured state. This is directly analogous to the busy server in a queuing system. Of interest is how the duration of these busy periods vary with the implementation of a full- or half-duplex channel, the data transmission rate, packet size, and bit error rate.

The duration of user sessions directly impacts the number of users who may actually gain access to the net in a finite period of time. So, to a large extent the design parameters and questions enumerated for channel access also apply to the session duration problem. First and foremost is the decision to implement a full- or half-duplex net. In this context, engineers may then determine the direction to pursue regarding the synchronization process, bit rate and maximum packet length.

3.2.1 Duplex and Half-Duplex

Nominally, the sequence of events required to ensure a successfully completed half-duplex session is longer and more complicated than that for the same session on a full-duplex circuit. The key concept is that all events in a half-duplex channel are serially arranged. In a full-duplex net, some of these events occur simultaneously. Recall, however, that the transmitted power is spread over half of the band-width in a full-duplex channel, because the net allows for simultaneous communications in both directions. Contending users' transmissions may contribute greater noise to the environment than in a half-duplex net. Sessions may be conducted over a circuit subject to the higher bit error rate resulting from the potentially noisier environment. A question of interest is whether the implementation of a full-duplex network results in significantly shorter sessions regardless of the bit error rate.

3.2.2 Packet Length

Shorter packet lengths mean shorter sessions. Not only because the subscriber uplinks shorter messages, but all those stored for download are also shorter. Furthermore, shorter packets have a smaller probability of experiencing bit error than longer packets. This further eliminates the occurrence of sessions extended due to retransmitting lost data.

It is preferred to limit communications as little as possible, however. Placing too restrictive a constraint on packet length may impede a subscriber's utilization of the network. The engineers' objective is to allow for as high an upper bound on packet size as feasible while preserving the goal of minimizing session duration.

3.2.3 Spreading Sequence and Data Transfer Rates

Pertaining to session duration, these two design issues have a similar impact as discussed regarding channel access. Engineers are interested in determining how much the synchronization sequence must be shortened to significantly decrease the duration of a session. Again, it is somewhat apparent that doubling the data transfer rate will shorten a typical session. But, is this enough to overcome any loss in the margin of the link analysis? PANSAT's designers want to know if doubling the data rate results in significantly shorter sessions despite a higher bit error rate.

3.3 SCU Storage Requirements

The driving force behind PANSAT is its utility as an experimental platform. Engineers are concentrating on designing the spacecraft to establish the operability of this low-orbit communications network. Still, the storage of packets poses a design problem. SCU memory used for packet storage varies with the number of users and packet size.

It is important to consider that over the course of its lifetime, interest in PANSAT may grow to the extent that design limits on storage capacity may be flexed to the limit. For example, the satellite's software designers claim that the SCU's operating system will be altered significantly if the storage size exceed 500 kilobytes of data. Is this much space required or will this level be exceeded? Design matters here must show robustness over a spectrum of resolution, introducing greater variability in storage levels.

To determine what storage capacity to design for the spacecraft control unit's memory, engineers are interested not only in a specific measure such as the average or maximum level of storage used. Designing based upon an average does not reflect the range over which the running total walks. As packets are received and are downloaded, observing fluctuations in storage level will generate a mean value; the estimator gives no information on the proportion of observations which exceed this point, or the frequency with which they do so. Similarly, measuring the high-water mark, the maximum level reached by the SCU's memory, provides no description regarding how often this occurs. In any case, the measure of interest is the amount of memory in use at the time of the next packet's receipt.

Decision makers want to know the distribution of memory used. From a design standpoint, the principle factors in this density are the number of users allowed to participate in the net and determining the maximum packet length. These two parameters may be readily used in an analytical solution. However, the frequency with which SCU storage varies over a range of values is very much scenario dependent. For a given number of subscribers and packet size selection, designers want the SCU to accommodate a reasonable proportion of the total number of messages over a number of diverse scenarios.

3.4 How a Model Helps

In summary, a partial listing of influential design parameters follows:

- data transmission rate;
- user population and population density distribution;
- bit error rate;
- minimum, maximum and fixed packet length;
- synchronization attempt discipline;
- SCU storage capacity;
- full- or half-duplex communications.

Complications arise in trying to assess the interdependency of these design factors and their effect on the network. With this level of complexity what kind of numbers may be generated by a model? How will the observed simulated activity differ from specified probability models applied to each question of interest?

For the model's results to be credible, a high degree of resolution must be incorporated in the representation of network communications. In analyzing

the system and creating the objects which will emulate the flow of activity, certain characteristics emerge which suggest a common-sense approach or analytic probability model both of which may resolve a design issue more quickly and as effectively as a simulation run. The model's three levels, common sense, analytic and simulation, complement each other.

For each measure of performance, obtaining expected values may be consistent between the probability models and the simulation. However, as a descriptor of performance, the average value falls far short. In essence, the merit behind thorough representation of network activity is the insight into system variability afforded the decision maker. This is where the anticipated divergence is between computer runs and paper or blackboard results. In any event, the working model and analytical results will still pave the way for several major design decisions which may be further substantiated by a high-resolution simulation.

4 INTRODUCTION TO PACSIM DESIGN

This communications network may be modeled with a great deal of fidelity as a complicated stochastic model implemented in a simulation. The simulation is written in reusable, object-oriented code so that other design issues may be explored experimentally and so that new studies can be pursued as scenarios are developed by PANSAT's sponsors. Accompanying probability models are incorporated to validate the simulation.

4.1 Model Objects

In this section, we describe the objects used as building blocks in building PACSIM. Each of the objects is fully reusable, modifyable, and extendable.

4.1.1 Users

The network is activated by users attempting to capture the satellite. They are uniquely identified by the following attributes:

- callsign, for addressing purposes;
- synchronization attempt process, dictating the interattempt waiting period and number of retries;
- location, either isolated or collocated in a population center.

The model accounts for these essential characteristics as well as the user's network activities by creating each user as an autonomous object. In this regard, operations of these users are encapsulated in a series of object methods which define their impact on the network.

Channel access, packet flow and message handling, are significantly affected by subscriber actions. Successful synchronization causes the SCU to become busy and conduct a session. Subscribers generate messages. Message size and frequency of generation affect both the duration of the session as well as storage by the SCU. Transmission of the packet requires the passage of time. This influences the session duration and in turn affects other subscribers' ability to access the net. A user's traffic receipt process has several ramifications across the network. Successful receipt of packets from the spacecraft completes the end-to-end packet flow and also allows the SCU to free storage space for future use. If packet receipt is obstructed by bit error, session durations become extended and endpoint-to- endpoint communications flow may be disrupted. Finally, the procedures of acknowledging and requesting disconnect prompt the end of a session and close the loop on packet flow. These essential activities were selected for inclusion in the model because they:

1. fully describe the network activity of the users;
2. affect the state of the network in all areas, channel access, communications flow and packet handling;
3. preserve a high level of resolution in the model by reflecting actual network operation.

4.1.2 Spacecraft Control Unit

The SCU object contains two sub-objects, the user interface and message storage area. These two entities have attributes and procedures which allow for fully analyzing the SCU and its role in the network.

The message processor is the user interface. As discussed in the network description, it drives the session with the active user. When the SCU object conducts a session, it waits for packet transmission by the active user, commands all packet storage and retrieval, manages any acknowledgement procedures and terminates the session. In short, it completely emulates PANSAT's network communications.

Similarly, the storage object within the spacecraft module mimics all salient features of PANSAT's SCU storage. Once commanded by the control unit, the storage manipulates the messages and identifies all

packet addressees by callsign. The amount of packet storage space utilized is readily monitored; the value may be accessed at any time. These features fundamentally cover the activity of the storage unit in network operations.

4.1.3 Bit Packet

Just as in the actual network, the bit packet is generated, transmitted to the spacecraft, received, stored and retrieved and downlinked to the destination. It is subject to bit error at ends of the communications flow. Once the packet is routed to storage, space is allocated; upon forwarding, space is deallocated.

All network operations which influence the movement of bit packets are reflected with a high degree of fidelity in the model. This is essential to the model's utility in decision making. Properties of the model which prove crucial to its ability to properly imitate the network are discussed in the next section; these are followed by a listing of network features implemented in PACSIM.

4.2 Model Fidelity: Including Details

To create a credible model, a number of concepts relevant to network analysis must be incorporated in the simulation. Simulation is a unique tool, allowing the decision-maker to view the model network operation under varying conditions. The model includes flexibility in a number of areas which impact the questions of channel access, packet storage and endpoint-to- endpoint connectivity. They include:

1. capturing the SCU – PACSIM accurately emulates the processes of attempting channel access, synchronizing and conducting retires for channel access;
2. packet generating environment – the model incorporates any number of scenarios for imitating a user's need to generate a message, determining message size, and identifying to whom the message is addressed;
3. occurrence of bit error – ramifications of this event are fully incorporated in the simulation;
4. geographically sensible population density distribution – because the user-base is unknown, network performance may be observed using any number of scenarios distributing users throughout PANSAT's orbit;
5. dynamics of a low-orbit spacecraft – frequency of spacecraft views and the durations of these windows reflect the framework expected in the network.

A more verbose account of the details included in the model can be found in Gottfried [5].

5 PACSIM VERIFICATION AND MODEL VALIDATION

Before basing any design decisions on PACSIM, we must take two quality-assurance steps with the model. We must

- verify the simulation – ensuring that the program operates correctly;
- validate the model – ensuring that the representation of the PANSAT system is accurate.

Verification was performed by testing each of the forty (40) distinct MODSIM modules of PACSIM, testing bottom-up as is standard when programming objects. Validation in its purest form, where data from the real system is matched with simulated data, was impossible because the real system or anything like it does not yet exist. Our approach was to validate by

1. manipulating data and simplifying the model so it essentially matched an tractable model like an M/G/1 queue;
2. comparing the simplified model results to those obtained by solving analytical equations.

We pursued this for three different simplified models:

1. Making the satellite footprint stationary, with user sessions taking no time to transpire and message destinations equiprobable among all users, the number of messages in storage at any time is a tractable Markov chain.
2. Using full duplex communications with no bit errors, each session is the convolution of a set of random variables with known distributions.
3. Allocating user access attempts as a Poisson process, ensuring fully reliable synchronization on the first try and preserving a view of the satellite for the entire run, channel access of each user follows the same distribution as admission to an M/G/1 system.

In each case, PACSIM matched theoretical models very well.

As a final validation step, we analyzed several design decisions about which the engineers had strong intuition:

1. Quantity of SCU storage used with a 1000 and 2000 byte maximum message length increases with the number and population density of subscribers.

2. Session duration decreases when shifting from half- to full- duplex or when increasing data transfer rate from 1200 baud to 2400 baud communications.

3. Channel access is improved when data transfer rates increase.

6 PACSIM RESULTS

PACSIM ran nearly 200 different scenarios to explore the design questions enumerated in Section II. All observations were recorded during steady state in order to evaluate the system's long-run behavior. Session duration and channel access observations were measured at the completion of transactions while SCU memory observations were taken when messages were being received by the SCU. The storage, session duration and access issues listed in the preceding section follow.

6.1 SCU Storage

PANSAT's designers sensed that an increased number of users and larger maximum message sizes require a greater quantity of SCU storage. PACSIM confirmed these ideas and provided further insight into this design issue. Figure 5 summarizes the following conclusions:

1. in a heavy message traffic scenario and 2000 byte maximum message length, the initial 500 kilobyte memory threshhold set by the communications designers becomes an issue when there are more than 300 users in the network; the envelope is pushed further out for the 1000 byte message limit;

2. variability in the quantity of memory used increases with increased user population density; all experiments used four population centers – increased channel contention is associated with a decreased rate of access, creating an opportunity for messages to accumulate in storage;

3. for the 2000 byte maximum message length, memory usage increases superlinearly with the number of users; this corroborates the idea that channel access is linked to SCU memory usage.

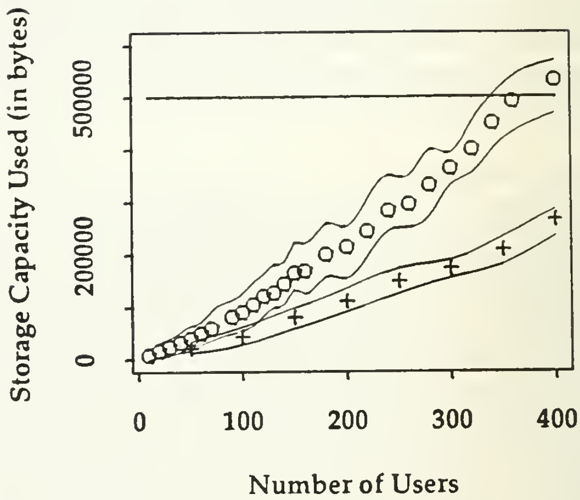

Figure 5: Average quantity of SCU storage used. Circles correspond to 2000 byte maximum message length, while plusses show the results for 1000 byte messages. The surrounding curves show the 95% confidence interval widths for each point. The horizontal line shows the current storage maximum of 500 kilobytes in the SCU.

6.2 Session Duration

Efforts to shorten session lengths and increase communications duty cycles include shifting from a half- to full-duplex circuit and increasing bit rate. These sessions were run under the heavy message generating environment, in which users uploaded a message and, on average, downloaded one message. Figure 6 summarizes the following observations.

- full-duplex is very robust against increases in bit error rate;

- full-duplex sessions have a significantly smaller rate of increase in duration as bit error increases;

- due to a greater number of prolonged sessions, half-duplex experiences more variability as bit error increases;

- on average, full-duplex sessions are shorter, even at bit error rates exceeding three times that of half-duplex sessions.

- 2400 baud communications are even more robust against bit error than the 1200 baud data transmission rate;

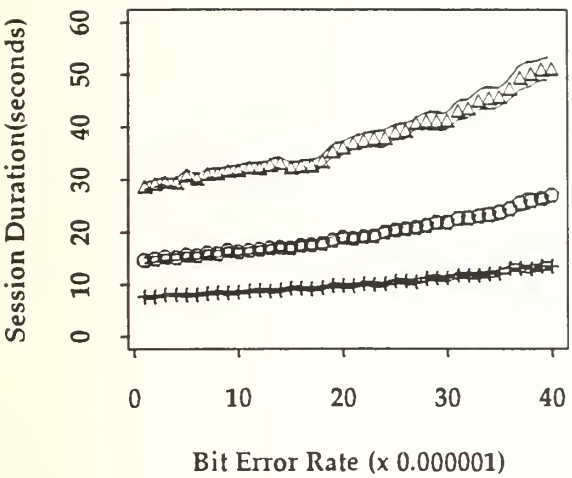

Figure 6: Session duration versus bit error rate. The circles represent values for the full-duplex system, while the triangles represent half-duplex sessions, both for 1200 BPS. The plusses are data for full-duplex, 2400 BPS. The surrounding curves show the 95% confidence interval widths for each point.

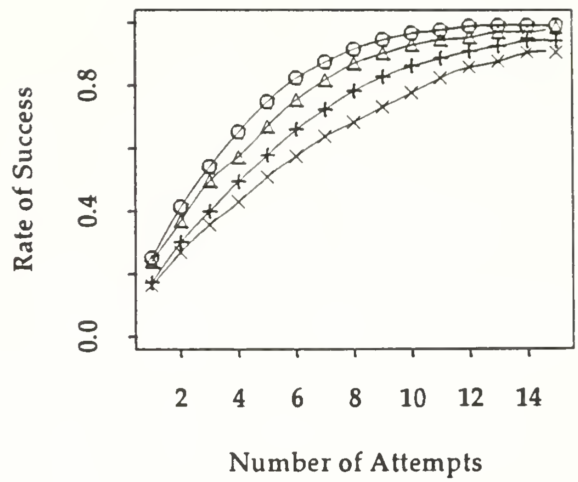

Figure 7: Proportion of user attempts resulting in successful access of the SCU. The circles represent data points for 2400 BPS for bit error probabilities of 10^{-6} and 3×10^{-6}. The plusses and crosses are for 1200 BPS, bit error probabilities of 10^{-6} and 3×10^{-6}.

- increases in session duration are almost insignificant even at a bit error rate four times the bit error rate assumed in the link analysis;
- if the assumed link margin remains constant, the 2400 baud circuit outperforms the 1200 baud circuit at a signal to noise ratio which allows for quadruple the bit error rate.

6.3 Channel Access

As discussed, design decisions with respect to channel access revolve around the following issues:

- the likelihood that the Naval Postgraduate School will be able to access the SCU unimpeded;
- reduction in the number of attempts users make to access the spacecraft – with each try, more noise is created on the channel;
- what proportion of users can expect to access the spacecraft.

The engineers can improve overall channel access by shortening the duration of sessions. We ran PACSIM scenarios measuring channel access against changes in data rate. Figure 7 depicts the following points:

1. even at three times the bit error rate, channel access is significantly better in a 2400 baud circuit than one at 1200 bits per second;
2. at 2400 bits per second, overall channel access is 90 percent successful after eight attempts, while this level of success is attained after 14 tries in a 1200 baud net.

In addition we ran the following experiments:

1. measuring session durations with maximum message lengths of 1000, 2000 and 4000 bytes;
2. measuring session durations with the maximum number of synchronization sequence repetitions set at 128, 96 and 64 iterations;
3. comparing the rate of channel access while varying exponential backoff rates of 15, 30, 45 and 60 seconds;
4. identifying the rate of channel access while varying the maximum number of synchronization sequence iterations among 128, 96 and 64.

7 SUMMARY AND CONCLUSION

The decision makers are engineers with vast experience in space systems. Conclusions based on their intuition benefit from a working knowledge of

PANSAT's design issues. PACSIM's first step toward credibility was confirmation of the intuitive assessments:

1. a larger number of users leads to increased SCU memory requirements;
2. shifting from half-duplex to full-duplex or increasing the data transfer rate shortens sessions.

The next step was to provide an understanding of variability in design issues:

1. increased contention for the SCU yields increased frequency of extended sessions and larger swings in SCU memory used;
2. longer messages not only increase session durations but experience greater variability in session duration due to retransmission requirements.

This also allowed for discussions of significant differences between performance measures for different designs. Finally, PACSIM has brought to light several issues which make it a useful tool for satellite network design and analysis:

- Multiple attempts are required to obtain a reasonable chance of accessing the SCU. If the Naval Postgraduate School ground control center wants to be assured of access, some other means, such as secondary frequency or coded access, will need to be implemented.
- Network design issues and factors are very densely interrelated. Session duration was expected to have a strong effect on all areas of communications flow. However, in a heavy traffic scenario, it has emerged over the course of experiments to be the driving force behind channel access and SCU utilization.
- Efforts to shorten session duration by improvements in effective data rate make the network more robust against bit error.

At this point in PANSAT's development, the insight into and measurements of characteristic network performance under various design decisions and operating scenarios cannot be obtained anywhere else. PACSIM is not a panacea, however. Numerical results are very much driven by scenario and cannot be extrapolated or predicted when applied to others. This is the reason for providing multiple levels of resolution to the model. Unfortunately, there is no way to ascertain whether enough reality has been implemented or if the correct scenarios have been incorporated in the simulation. It is difficult enough to forecast the utilization of a system which has been operated previously. PANSAT will be the first of its kind. PACSIM provides a structure to analyzing the potential operating environment of PANSAT and a means of assessing improvements and degradations based upon design decisions. This tool was not previously available to the design team. A model is a foundation and sets the tone for building a successful network. PACSIM provides a window on previously unforeseen communications and allows the engineers to make more informed decisions.

References

[1] Brachman, B. J., and Chanson, S. T. 1988. Fragmentation in Store- and-Forward Message Transfer. *IEEE Transactions in Communications* **26**(7), pp. 18-27.

[2] Bratley, P., Fox, B. L., and Schrage, L. E. 1987. *A Guide to Simulation*, 2d ed. New York: Springer-Verlag.

[3] Chapin, A.L. 1983. Connections and Connectionless Data Transmission. *Proceedings of the IEEE* **71**(12), pp. 1365-1371.

[4] Fossett, C. A., *et al.* 1991. An Assessment Procedure for Simulation Models: A Case Study. *Operations Research* **39**(5) pp. 710-723.

[5] Gottfried, R. 1992. *PACSIM: Using Simulation in Designing a Satellite Communications Network.* Master's Thesis, Department of Operations Research, Naval Postgraduate School, Monterey, California.

[6] Law, A. M., and Kelton, W. D. 1991. *Simulation Modeling and Analysis*, 2d ed. New York: McGraw-Hill.

[7] Morgan, W. L., and Gordon, G. D. 1989. *Communications Satellite Handbook.* New York: John Wiley and Sons.

[8] Panholzer, R. 1991. *Proposal for Research: Petite Amateur Navy Satellite (PANSAT).* Proposal for Research, Naval Postgraduate School, Monterey, California.

[9] Pursley, M. B. 1987. The role of Spread Spectrum in Packet Radio Networks. *Proceedings of the IEEE* **75**(1), pp. 116- 134.

[10] Tanenbaum, A. S. 1988. *Computer Networks*, 2d ed. New York: Prentice Hall.

[11] Woerner, B. D., and Stark, W. E. 1991. Packet Error Probability of DS/SSMA Communications with Convolutional Codes. *MILCOM '91 Proceedings*, pp. 127-131.

DISTRIBUTION LIST

1.	Defense Technical Information Center Cameron Station Alexandria, VA 22304	2
2.	Library, Code 55 Naval Postgraduate School Monterey, CA 93943-5002	2
3.	Research Office, Code 08 Naval Postgraduate School Monterey, CA 93943-5000	1
4.	Operations Research Department Naval Postgraduate School Monterey, CA 93943-5000	1
5.	Professor Michael Bailey Code OR/Ba Operations Research Department Naval Postgraduate School Monterey, CA 93943-5000	20
6.	Professor Rudolf Panholzer Code EC/Pz Electrical Computer Engineering Naval Postgraduate School Monterey, CA 93943-5000	10
7.	Capt. E. Enterline Director Navy Space Systems Operations Office of the Chief of Naval Operations (N63 Washington, D.C. 23060-2000	2
8.	ADM H. Browne Commander Naval Space Command Dahlgren, VA 22443-5170	2
9.	ADM Jay W. Sprague Commander Space and Navy Warfare Systems Command SPAWAR-004 Washington, D.C. 20362	2
10.	Russell Gottfried United States Navy Department Head School Class 126 Newport, RI 02841	20